arXiv:1607.03366v1 [cs.HC] 12 Jul 2016

Human-Planned Robotic Grasp Ranges: Capture and Validation

Brendon John[1], Jackson Carter[2], Javier Ruiz[3], Sai Krishna Allani[2], Saurabh Dixit[2], Cindy M. Grimm[2], and Ravi Balasubramanian[2]

***Abstract*— Leveraging human grasping skills to teach a robot to perform a manipulation task is appealing, but there are several limitations to this approach: time-inefficient data capture procedures, limited generalization of the data to other grasps and objects, and inability to use that data to learn more about how humans perform and evaluate grasps. This paper presents a data capture protocol that partially addresses these deficiencies by asking participants to specify *ranges* over which a grasp is valid. The protocol is verified both qualitatively through online survey questions (where 95.38% of within-range grasps are identified correctly with the nearest extreme grasp) and quantitatively by showing that there is small variation in grasps ranges from different participants as measured by joint angles, contact points, and position. We demonstrate that these grasp ranges are valid through testing on a physical robot (93.75% of grasps interpolated from grasp ranges are successful).**

I. INTRODUCTION

Humans are excellent at physical manipulation, robots less so. Thus, it is appealing to use humans to "teach" robots how to manipulate objects. This difference in abilities is not just a result of the human and the robot using different manipulators; that is, the human hand versus a robot hand. There is a difference in the strategies humans use when compared with the strategies used by current robotic grasp planning. For example, the work in [1] showed that humans use even robotic manipulators better than current automatic algorithms for unstructured grasping tasks. Unfortunately, it is unclear what strategies humans use to choose a grasp. Furthermore, elucidating those strategies is challenging. Thus, most methods used to train robots rely on some form of human demonstration of *specific grasp examples for specific tasks* using a physical robot [2]. Unfortunately, this method of collecting human input only results in a small number of specific grasps for specific objects. It is not clear how to generalize this data or make it robust to small perturbations in robot posture or object location.

In this paper, we outline a novel protocol for collecting human-generated grasp examples that capture not just a single grasp, but an acceptable *variation* of that grasp from many participants. This protocol results in a more time-efficient capture of example grasps and data that is easier to generalize. Specifically, the protocol: 1) Captures *regions* of good grasps *and* bad grasps, instead of single ("optimal" or first-choice) grasp instances. 2) Captures additional human-centric information to elucidate: *how* the participant arrived at a specific grasp, how that grasp compares to a grasp that the participant would do with their own hand, and what the participant looked at when performing the grasping task. 3) Supports capture of manipulation tasks where one robot hand can only perform part of the manipulation, such as picking a small object off the table or re-positioning a grip. Post-protocol, we also use a survey-based approach for obtaining information on the similarity of grasps collected in order to better group them.

Capturing a grasp *region* instead of a single grasp, and capturing many bad grasps, is useful for two reasons. First, sophisticated machine learning algorithms can use this data to better define regions of valid inference, such as the *boundaries* of where a successful grasp will fail. Increasing the number of negative training examples is also important to improve machine learning results (for example, where not to focus grasp search). Second, this provides insight into how humans group grasps. This information can be used in human-robot interaction applications to search for grasp metrics that match those groupings [3].

Identification of similar grasps across people and objects makes it easier to generate robust grasps by providing a partitioning of the data set into grasps where simple linear interpolation makes sense. In our experiments, the first few grasps the participants provided tended to be similar to the ones given by other participants, providing evidence that these grasps are reliable in the view of the participants.

We use a mix of a think-aloud protocol, prompting questions, and eye-tracking to capture human thought processes and visual attention during the grasping tasks. We also ask the participants to perform the grasp with their own hand as well as with the robotic hand, and explain the differences. This captures both high-level cognitive processes (what the participants *think* they're doing) and low-level actions (actual physical and perceptual actions). Although we do not discuss analysis of this data in detail in this paper, preliminary analysis shows that this information is useful for identifying good views of objects and an initial (partial) mapping of existing grasp metrics to human considerations. [4]

The proposed protocol supports capturing a variety of grasping and manipulation tasks, such as picking an object up from the table, receiving the object from another hand, and manipulating the object. To test our protocol, we chose one structured task for each object (pick the object up from the table and place it on a box), and one more natural task suit-

*Supported in part by NSF Grant REU Site: Robots in the Real World, CNS 1359480

[1]Rochester Institute of Technology, bmj8778@rit.edu

[2]Oregon State University, allanis, carterj, dixits, Cindy.Grimm, Ravi.Balasubramanian, @OregonState.edu

[3]UC Santa Cruz jamaruiz@ucsc.edu

able for that object (such as pouring for the water pitcher). The latter was deliberately free-form — the participants were free to place the object in the robot hand or pick it up from the table. To support this more free-form interaction we used a Kinect™ sensor to track the object's location relative to the hand (in addition to tracking the joint angles of the robot hand and arm).

To validate the grasp ranges collected, we focus on the following questions: 1) How do we identify and verify similarity of grasps across participants? 2) What is the variation in the range of grasps that participants prefer for an object? 3) How do grasps sampled from the participants's grasp ranges perform in practice?

II. RELATED WORK

There has been significant progress in the domain of robotic grasping and manipulation both in terms of hardware [5], [6], [7] and software development [8], [9], [10], [11]. However, prior work has shown that even in a laboratory environment with almost perfect information for grasp planning, robotic grasping performance only succeeds about 75% of the time; that is, one in four grasps fail [1]. The primary reason for this poor performance is that robot grasps are not robust enough; that is, small differences in object shape or object position cause the object to, say, slip out during the grasping process. There has been significant effort to address these issues using physics-based heuristics and brute-force search algorithms to find more robust grasps [12], [13]. However, a big hurdle that these methods face is that grasping tasks exist in a continuous space with significant uncertainty (for example, arising from unknown friction and compliance in the gripper fingertips) and noise (for example, arising from the actuator position errors or warped sensory perception). Thus, even if perfect discrete grasps are created, a millimeter or two of positioning error will cause the grasp to fail. This is precisely the reason why this work collects grasp *ranges* from humans and identifies grasps that participants rank as similar. This provides evidence that the grasp is robust to small variations and grasps do not live on discrete islands.

Prior work has also explored "learning from demonstration (LfD)", where humans teach robots to advance robot performance [14], [15], [16], [17]. The key idea is to identify generalized control policies based on specific examples. However, these approaches for gathering data are time-inefficient and require vast amounts of data to learn policies [18]. The approach taken in this work of collecting grasp ranges fits well with existing LfD techniques [19], since it taps into human intuition for generalizing grasps around optimal grasps.

III. STUDY PROTOCOL AND METHODOLOGY

In this section, we describe the human-subject experiment protocol in terms of the tasks, phases, specific queries, data captured, protocol management, and participants. At a high level, participants were asked to perform a specific manipulation task with a given object. They accomplished this task both with their own hand and with the robotic hand, but we used only the data from the trials with the robot hand. The participants were asked to show as many "good" and "bad" grasps as they could. In addition to performing the manipulation, they were asked to specify the valid range for each unique grasp (if such a range existed). We captured four data streams: the participant's eye gaze, audio, arm and hand posture, and a 4D video (color plus depth) (see Table I), but use only the robot arm and hand position and 4D video in this paper.

TABLE I
DATA CAPTURE METHODS AND USAGE

Data	Usage
Gaze Data	Identify which regions of the scene are fixated on most while grasping.
Eye Tracker Scene Video	Used to get perspective view of subject, provides snapshots of specified grasps.
Kinect Depth Data	Object tracking and grasp reconstruction.
Barrett Hand and Arm Data	Joint angles and positions used for recreating grasps performed in the study.
Human Specified Ranges	Ranges can be utilized to improve machine learning models, and tested using an automated system.
Good and Bad Trials	Bad grasp trials provides samples of grasps to avoid when working with models or simulations.

A. Objects and tasks

The objects used in this protocol are shown in Figure 1. For each object, the participants were asked to perform two tasks. The first was to pick up the object from the table and place it on a nearby box. The second task was object-specific (see Table II). These object-specific tasks are tasks or actions that are commonly associated with each object, such as throwing a ball, squeezing a trigger, or handing the object to someone.

For some of the tasks (such as picking up the snowman from the table) the robotic hand was not physically capable of performing the task with the object lying on the table. In this case, the participants were allowed to pick up the object and put it in the robot's hand. However, the participants were required to actually perform the task with the robotic arm, such as grasping the object with the robot hand and then moving the arm to perform the task. This ensured that the given grasps were successful, even if placing the object into the grasps required assistance.

B. Phases

Our study protocol is designed to capture both human grasping and human-planned robotic grasping. To do this, the study features a training phase and two distinct capture phases.

In the training phase, which preceded data capture, participants were asked to familiarize themselves with the robot hand by moving it around and adjusting the fingers, specifically showing them how to change the spread of the fingers. Although the arm was gravity compensated, it did not always maintain its position when the hand was opened and closed due to the hand's shifting center of mass, so

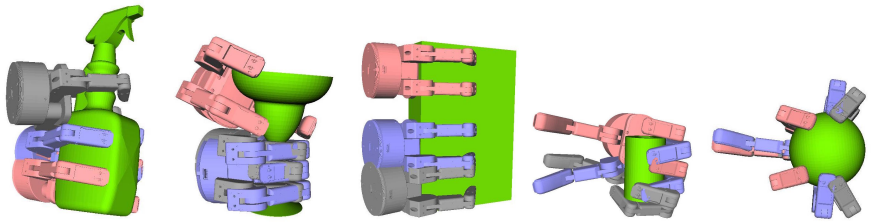

Fig. 2. 1-5: Five objects with the grasp range given by one participant. The violet hand is the original grasp, the gray and red show the first and second extreme grasp locations.

Fig. 1. Objects used for the study.

participants were also given instructions to ask for help in supporting the hand if needed.

In the first capture phase, the participants use their own hands to grab an object, while in the second, the participant physically positions the robotic arm and hand to grasp the object. The order of the two phases was randomized for each object.

For the human-hand grasping phase, participants were asked to use only their thumb and first two fingers to mimic the three fingers of the robotic hand. This paper does not the human-hand grasping data, which will be analyzed in future work.

The robotic grasping phase is further split into three phases: Specifying a pick-up grasp, performing the manipulation task, and specifying "bad" grasps. For every grasp we also asked the participants to specify a range for that grasp (if the participant felt there was one) before moving on.

Grasp ranges: For each specified grasp (both good and bad), the participant was also asked to provide a range along the object for which that particular grasp could still be applied (e.g. along the side of a box — see Figure 2)) . Participants were also asked if there were any rotational or symmetric ranges for their grasps. In the good grasp phase, these ranges are the areas on the object that the participant believes the grasp would still be successful. In the bad grasp phase this range represents the region that a particular grasp would still be ineffective. Not all grasps had ranges, for example, holding the spray bottle to spray.

Bad grasps: Participants were asked to "teach" the robot how *not* to pick up the object by demonstrating bad grasps. They were allowed to specify as many grasps as they wanted, but were asked to demonstrate the grasp failing for each one[1].

Because of participant time constraints and the eventual fatigue induced by the procedure, not all subjects were able to complete all phases of the study. Occasionally, participants would complete the good grasp phase of the study, but not have time to complete the bad grasp phase. Thus, some objects do not yet have bad grasp data and associated statistics.

TABLE II

OBJECT-SPECIFIC TASKS AND NUMBER OF GRASPS CAPTURED (INCLUDING PICK-UP TASK).

Objects	**Natural Task**	**Grasps**	
		Good	**Bad**
Water Pitcher	Pour water out of pitcher	11	1
Spray Bottle	Pull trigger to spray	14	22
Margarita Glass	Drink out of glass	14	14
Cereal Box	Pour cereal out of box	12	18
Cracker Box	Pour crackers out of box	15	7
Television Remote	Press power button on remote	11	14
Toy Plane	Pretend to fly plane around	13	19
Food Clip	Open clip (to close bag)	10	3
Soap Dispenser	Press down on nozzle to dispense soap	10	5
Foam Cylinder	Throw object overhand	16	15
Bison Plush Toy*	Hand toy to someone	5	0
Plush Ball	Throw ball underhand	19	10
Thunder Stick*	Hit something with it	10	0
Sock Doll	Hand doll to someone	16	15
Decorative Cord	Hang cord by its metal ring	5	6
Tape Roll	Support tape roll for ripping tape off	11	4
Total		192	153
Mean		11.29	9

Note: Objects with * do not have any bad grasps

C. Prompts and think-aloud

The subjects were asked to think out loud as they performed the study to provide insight into what they were thinking of while performing the grasping tasks. For the pick up the object task, the participants were asked to move the object using the robotic hand after finalizing the grasp. For the other tasks, they were not required to perform the task, but simply needed to position the hand. They were given explicit permission to pick up the object, position it how they wanted, and to use their other hand if they needed two hands. Participants were always given the option of

[1]Some participants accidentally found good grasps this way.

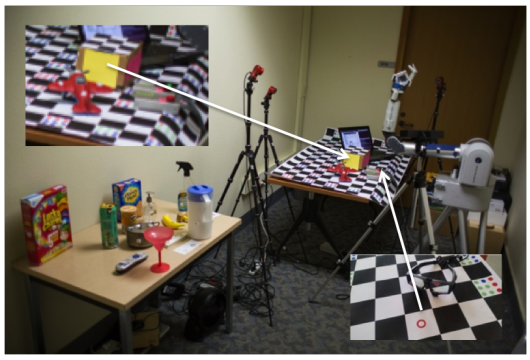

Fig. 3. Study set up. The table included a checkerboard pattern for further calibration. The red circle was used to calibrate the eye tracker. The box on the table was used in the object placement tasks. Participants were initially seated on the opposite side of the table from the Kinect sensor; they were not required to remain seated.

specifying another grasp for that task (if they could think of one).

We had one further prompt, asked at the end of every good grasp: "Is this grasp exactly what you wanted? Or are the finger placements slightly different that what you were intending? (How so?)". This prompt is aimed at disambiguating how much the robotic hand limitations affected the participant's grasp choice.

D. Data capture: equipment and procedure

The equipment used for this study included a pair of SMI Eye Tracking Glasses 2.0 to collect eye-gaze data, a Microsoft Kinect v2 to collect depth data, and a Barrett Whole Arm Manipulator (WAM) with BH280 BarrettHand to perform the robotic grasping. We also included spatial calibration patterns for the table and box (see Figure 3) to ensure calibration between video data sources. Data captured, and its uses, are summarized in Table I.

Eye tracking: The SMI glasses record both where the user is looking and what they are looking at. The data is recorded as a 960x720 video stream at 30 Hz, plus an eye gaze location for each video frame (as x, y image coordinates). The eye gaze data also includes other information such as pupil diameter, fixations, and saccades. The eye tracker has to be fit to the person's head (similar to goggles) using two nose pieces and calibrated to their eyes. To perform the calibration the participant was asked to sit down in front of the table and fixate on a red dot on the table (see Figure 3). This one point calibration was performed using the SMI software. We checked the calibration at the end of each grasp trial by having the participant focus on the red dot again. The eye-gaze tracking data is not used in this paper.

Arm and hand tracking: We used a Barrett WAM and BarrettHand (BH-280) in the study. The arm is backdrivable and gravity compensated; that is, the arm location can be physically adjusted with ease. However, the BarrettHand's fingers cannot be physically adjusted from external forces (only through its motors). We used a physical set of three sliders to control how much each finger was closed, and a knob to control the spread of the fingers. Note that the two joints of the finger are controlled with one actuator.

The Barrett WAM itself features an onboard Linux PC which both compensates for gravitational effects on the arm linkage and continuously streams joint positions to the main Linux PC via ethernet. The onboard PC includes a timestamp with each joint position measurement, so care was taken to synchronize the onboard PC's clock with the main Linux PC's clock using the program chrony.

The gravity compensation function was calibrated with the hand open to reduce the need to support the arm when the hand changed configuration; however, it was still sometimes necessary to externally stabilize the hand when adjusting the fingers. We provided explicit instructions to the participants to ask one of the study team members to help if this happened.

Kinect 3D scene data The Microsoft Kinect provides a 960x540 RGB image with depth data for each pixel at 30fps. The depth is from time-of-flight, so is (approximately) correct in real-world terms. The depth data, along with an estimate for the camera's orientation with respect to the scene, was used to construct a 3D point cloud of the scene. We use the point cloud to find the object's location with respect to the hand (see Section V-A).

Audio: The eye-tracker records audio with the video.

Temporal data stream alignment: The Robot Operating System (ROS) was used to integrate and store data feeds from the Barrett WAM, the BarrettHand, and the Kinect. RGB and depth images are recorded continuously during the grasping phases along with the robot's arm and hand joint positions when it is being manipulated. When participants finalize a grasp, a fixed-format annotation is (manually) added to the data stream along with a timestamp for future synchronization. Additionally, a more general text annotation system is used to flag grasps in order to make them easier to find in the video data. The eye-tracking system was temporally synced with the other data streams using a high-pitched beep.

E. Protocol management and flow

The study is designed to be run by two researchers; one for Eye tracker and other for Ubuntu ROS PC.

The average time for a data collection session was an hour and a half, covering two grasps each for three or four objects. The maximum time was capped at two hours due to eye strain generated by the eye tracking glasses, as well as general fatigue from performing the experiment. Although we limited a single session to two hours, several participants did two sessions. The single-session participants were asked to perform good grasps on two objects first, and then bad grasps with two objects. For two-session participants we did good grasps for as many objects as possible in the first session, and bad grasps in the second.

The general flow of the study can be seen in Figure 4 and is also outlined in the list below.

1) Participant enters room and signs consent form.
2) Brief training session with a test object.
3) Eye tracking calibration performed.

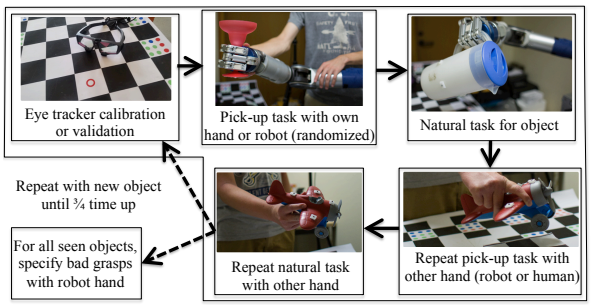

Fig. 4. Flow chart of study procedure.

4) Study trials explained to participant.
   a) Object placed on table, and participant told to use robot hand (group 1) or their own hand (group 2) to perform pick up task.
      i) Pick up task performed.
      ii) Grasp range specified
      iii) Repeat until no new grasps.
   b) Natural task explained
      i) Natural task performed.
      ii) Grasp range specified
      iii) Repeat until no new grasps.
   c) Object-tasks (a-b) repeated with human hand (group 1) or robot hand (group 2)
   d) Eye tracking recording stopped, and re-calibration if needed.
5) Repeat a-d with as many objects as possible (approximately 1.5 hours)
6) Bad grasp prompt
   a) Participant given first object from good grasp trials
      i) Bad grasp and range specified.
      ii) Repeat until no new grasps.
   b) Repeat a) with next object from good grasp trials
7) Eye tracking recording stopped, all other data collection ended.

F. Participants

We recruited 13 participants, ranging in age from 16 to late 50's, all with normal or corrected to normal with contacts vision (it is not possible to simultaneously wear regular eye glasses and the Eye-gaze ones). On average participants specified 2-4 good grasps per object, with a high of 8, and 3-5 bad grasps (see Table III). The largest number of grasps any participant specified for an object was 16 for the Foam Cylinder. The fewest number of grasps specified for an object was 4 for the Margarita Glass. Of all 294 grasps collected, 179 had extremes (therefore comprising a grasp range) while 115 did not.

IV. VERIFICATION OF PROTOCOL

In this section, we provide information on how we verified initial measures of grasp similarity across participants, how we verify and use the grasp ranges, and how we verify using physical shake tests and participant evaluation the quality of grasps generated from the grasp ranges.

Data collection of this sort is very time consuming. Furthermore, there is an inherent trade-off between data reliability and data variability. Specifically, we can collect a lot of (possibly redundant) data about one object manipulation task, or collect (possibly biased or missing) data for many object manipulation tasks. We have settled on the following protocol for determining how many participants is enough. We collect data from a minimum of two participants for each object. If the participants both specify one (or more) "similar" grasps (see Section IV-A) then we mark that object as complete and move to the next one. If the participants gave different grasps, then we continue until we see no new grasps. In practice, this did not happen.

The motivation behind this approach is as follows: 1) In our formative studies, participants were remarkably consistent in producing the same grasp and grasp range. While some participants were more creative (specifying many grasps), their first grasp was almost always the same as other participants. 2) Since we are explicitly collecting a *range* for each grasp, we do not need a lot of individual grasp examples. This range can then be verified with other, less time-intensive methods such as using crowd sourcing (showing images of example grasps within that range) or semi-automated physical shake testing [1].

A. Grasp similarity

We evaluate the similarity of grasps in three stages: 1) An initial qualitative assessment by the experimenters of the grasps provided for five objects to determine if the claims that participants provide similar grasps is true. 2) A subsequent quantitative assessment using joint angles, contact points, and positions. 3) Shake test of grasps grouped by similarity.

Stage one used images of grasps of five of the thirteen objects (cereal box, ball, chunk of foam, spray bottle, wine glass) to qualitatively segregate grasps into groups based on their similarity (see Figure 5). Stage two is quantitative: Although we do not have enough participants per object to make any statistically significant claims, these numbers form a base-line for quantifying the qualitative notion of similar grasps. All quantitative measures are normalized (zero is no variation, one is maximum variation). See Table IV. In stage three we performed shake tests of groups of similar grasps. We performed similarity analysis based on following metrics:
Joint angle: We measure the difference to the mean over the available range of the joint, averaged over all of the joints for the hand (finger spread, finger joint angles).
Contact points: We determine contact points at the fingertip, pad, and finger joints using OpenRAVE [20], specifically, we calculate intersections with the object and the hand after alignment (Section V-A). The measure we used was the average of all-pairs edit distance (number of changed contacts over number of contacts).
Positions: For this measure, we aligned the hand to the coordinate system of the object. For objects with symmetry

TABLE III

AVERAGE NUMBER OF EACH TYPE OF GRASP GIVEN PER PARTICIPANT PER OBJECT.

Type	Average number of grasps per participant												Mean
Good	4	6.25	3.25	4.5	4.8	7	3	3.5	3.75	2.5	2.5	5.67	8	4.5
Bad	0	7.5	4	5.75	3.2	0	3.75	4.75	4.5	1.5	2	2	0	3.0

TABLE IV

STANDARD DEVIATION BASED NUMERICAL EVALUATION OF SIMILARITY OF GRASPS (THREE PARTICIPANTS EACH).

Object	Grasp set	Type	Joint angle	Contact points	Positional distance
Spray	1	optimal	0.10	0.07	0.12
	2	optimal	0.08	0.00	0.03
Glass	1	optimal	0.08	0.13	0.12
	2	optimal	0.07	0.04	0.16
	2	extreme1	0.06	0.03	0.14
	3	optimal	0.10	0.15	0.21
Box	1	optimal	0.06	0.00	0.03
	1	extreme1	0.12	0.00	0.02
	2	optimal	0.06	0.03	0.16
	2	extreme1	0.06	0.00	0.23
Foam	1	optimal	0.08	0.00	0.34
	1	extreme1	0.09	0.02	0.38
	1	extreme2	0.09	0.00	0.36
	2	optimal	0.20	0.18	0.35
	2	extreme1	0.04	0.12	0.52
	2	extreme2	0.04	0.10	0.20
	3	optimal	0.03	0.00	0.05
Ball	1	optimal	0.05	0.09	0.20
	1	extreme1	0.08	0.10	0.29
	1	extreme2	0.10	0.07	0.22
	2	optimal	0.04	0.09	0.08
	2	extreme1	0.08	0.13	0.19
	2	extreme2	0.03	0.00	0.05
	3	optimal	0.03	0.03	0.12
	3	extreme1	0.24	0.15	0.26
	4	optimal	0.07	0.10	0.07
	4	extreme1	0.20	0.12	0.18
Mean	-	-	0.08	0.07	0.19

we used the palm's location to resolve the ambiguity (eg, the vector to the palm to align the glass' horizontal axis). We chose 14 positions (points) on the hand (three along each finger, five on the palm) and calculated their average distance from the mean, divided by the bounding box of the object.

B. Grasp ranges

We validated grasp ranges for the five objects both qualitatively (using an online survey) and quantitatively using a shake test (see following section). To generate intermediate grasps we employed the following algorithm implemented in OpenRAVE: Linearly interpolate the joint angles for the two (or more) grasps to be interpolated. Deal with any intersections by moving the palm, followed by the fingers, out of the surface. Then close the fingers until contact is made for those fingers that have contact points in one (or more) of the original grasps.

We validated the human-evaluated similarity of the interpolated grasps to the original grasps using a survey. Participants were shown three images: the original grasp, an interpolated grasp, and the corresponding extreme grasp. Participants were asked to pick which grasp (original or extreme) was most similar to the interpolated one. We tested two interpolated grasps; one that was 1/3 between the two and one that was 2/3. Figure 6, left, shows the results; in nearly all cases the survey participants correctly identified which grasp the interpolated one was most similar to (from a linear interpolation standpoint).

We also tested these interpolated grasps for effectiveness using a shake test (summarized in the following section; results of the shake test are in Table V).

C. Shake tests

We validated the original, extreme, and interpolated grasps using a shake test where the robot hand automatically picked up the object from the table and shook it (five times per test). The exception to this is the grasp with the hand under the ball; for this test we suspended the ball in the air using a string. We used OpenRAVE's inverse kinematics module to calculate the joint states for the robot arm to take the hand to the object. The robot's hand position itself was the one specified by the participant.

Table V summarizes the grasps and objects used in the shake test and the results (see also Figure 2, Right). The grasps were grouped by similarity across participants; different sets had different numbers of grasps within it (Count).

D. Grasp quality

In addition to validating the grasps using a shake test we qualitatively assessed grasp quality by directly asking the participants how they felt about the grasp they provided. Around half of the participants said at least one grasp was not quite what they wanted, particularly for more complex objects such as the plane. The major refrain was that the participants didn't like that the joints in the fingers couldn't be controlled individually (the Barrett fingers are underactuated so the slider bends the entire finger, not each joint independently). This was most noticeable in cases where the finger locks up due to collision — one part of the finger comes in contact and locks, while the tip or remaining part of the finger stops where it is and doesn't close all the way around the object. Other issues were the fingers being too thick, the hand too big, or the controls being too fidgety to achieve some of the more precise grasps the participants had intended to perform.

Only one participant commented that they might change their robotic hand grasp after doing the human hand task. We found no evidence for order effects in the data.

We also performed a post-capture on-line survey asking participants (not from the study that collected the grasps on the physical robot) to evaluate if the grasps (original, interpolated, and extreme) would work. For this survey, we used OpenRAVE renderings of the hand and object taken from the viewpoint (approximately) of the eye-gaze video. See Figure 6, right.

(a) Spray Bottle (b) Margarita Glass (c) Cereal Box (d) Chunk of Foam (e) Ball

Fig. 5. Images of similar grasps for three participants for five objects.

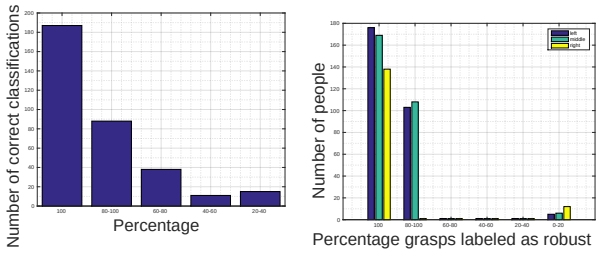

Fig. 6. Results of the post capture on-line survey (8-10 answers per grasp). Left: Distribution of intermediate grasps which were correctly classified; over half the grasps had 100% correct classification. Right: Distribution of grasps that participants said would work. Green is the intermediate grasp; blue and yellow are the optimal and extreme grasps.

TABLE V

SHAKE TEST RESULTS

Object	Set	Type	Count	Total	Successful
Spray Bottle	1	original	5(2+3)	25	24
	2	original	2(0+2)	10	10
	3	interpolated	2(2+0)	10	10
Wine Glass	1	original	3(0+3)	15	15
	2	original	6(3+3)	30	27
	3	original	3(2+1)	15	15
	4	interpolated	2(2+0)	10	5
Cereal Box	1	original	6(4+2)	30	30
	2	original	5(3+2)	25	25
	3	interpolated	2(2+0)	10	10
	4	interpolated	2(2+0)	10	10
chunk of foam	1	original	9(6+3)	45	44
	2	original	7(4+3)	35	33
	3	original	4(2+2)	20	20
	4	interpolated	2(2+0)	10	10
	5	interpolated	2(2+0)	10	10
Ball	1	original	7(4+3)	35	35
	2	original	8(5+3)	40	40
	3	original	5(3+2)	25	25
	4	original	5(3+2)	25	21
	5	interpolated	2(2+0)	10	10
	6	interpolated	2(2+0)	10	10
Total			91	455	434
			Overall grasps:		(95.38%)
			Interpolated grasps:		(93.75%)

V. DISCUSSION AND CONCLUSIONS

Protocol variations: Randomizing the order of human hand versus robotic one does not seem to produce any noticeable differences in the quality of the data. We have experimented with doing all human (or all robotic) first; however, this introduces considerable fatigue for the robotic hand phase, so we recommend interleaving them. Using two (or more) Kinect sensors would also improve the 3D tracking of the object.

Usefulness of data: Our on-going analysis indicates that there are important differences in the way participants approach placing the robotic hand versus the human one, both in what they gaze at and the types of grasps produced. Examples of *using* the captured data (eg analysis of what people were looking at or thinking, using machine learning on the captured grasp ranges) will be explored in future work. We will share the captured grasp data using RoboEarth [21].

Conclusion: We have presented a time-efficient protocol for capturing human-specified robot grasps for object-manipulation tasks that captures five different streams of data in the same amount of time. This protocol supports collection of *ranges* of grasps, potentially yielding more useful information for machine learning algorithms. The protocol also explicitly aims to capture the human reasoning behind the grasps through three mechanisms: Comparison with human hand grasping, think-aloud protocol, and eye-gaze tracking. We have verified that the extremes of the grasp

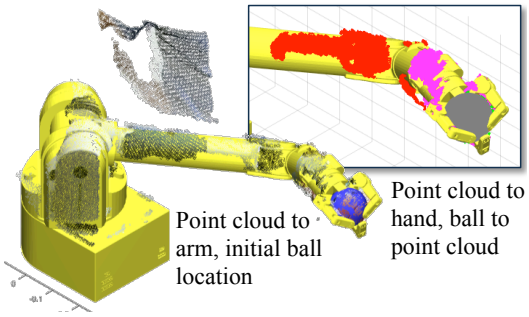

Fig. 7. Aligning the Kinect point cloud with the entire arm and hand. Upper right: Aligning the point cloud to the hand and the object to the point cloud by partitioning the point cloud (red arm, purple hand, green object).

ranges given by the participants do result in good grasps.

APPENDIX: 3D DATA CALIBRATION

A. Hand to Kinect sensor to object

Although we know the joint angles of the hand and arm, we do not know where the object is with respect to the hand. We used the Kinect point cloud to calculate an object transform in a three-step process: Generate a polygonal model of the arm from the joint data, align the point cloud to the arm, then align a polygonal model of the object to the point cloud (see Figure 7).

We used OpenRave [20] to produce a polygonal model of the hand and arm in the position specified by the joint angles. We hand-aligned this model to the Kinect data then used MATLAB's Procrustes alignment to refine this alignment matrix [2]

We used Iterative Closest Points (ICP) [22] to align the arm to the point cloud (trimmed to the arm area). We manually specified a starting transformation for the object then used ICP followed by Procrustes to alternate aligning the point cloud to the hand and the object to the point cloud.

ACKNOWLEDGMENT

Research funded in part by NSF grant CNS 1359480. We would like to thank the participants for their time and effort. We would also like to thank Dr. Reynold Bailey for the generous loan of his SMI glasses, without which this project would not have been possible

REFERENCES

[1] R. Balasubramanian, L. Xu, P. Brook, J. R. Smith, and Y. Matsuoka., "Physical human interactive guidance: A simple method to study human grasping." *IEEE Transactions on Robotics*, 2012, dOI: 10.1109/TRO.2012.2189498. (In press).

[2] S. Ekvall and D. Kragic, "Interactive grasp learning based on human demonstration," in *Robotics and Automation, 2004. Proceedings. ICRA'04. 2004 IEEE International Conference on*, vol. 4. IEEE, 2004, pp. 3519–3524.

[3] A. Billard, S. Calinon, R. Dillmann, and S. Schaal, "Robot programming by demonstration," in *Springer handbook of robotics*. Springer, 2008, pp. 1371–1394.

[4] M. S. Cindy Grimm*, Ravi Balasubramanian, "Visual cues used to evaluate grasps from images," in *(In Press) Conference: 2016 IEEE International Conference on Robotics and Automation*. IEEE, 2016.

[5] A. M. Dollar and R. D. Howe, "The highly adaptive SDM Hand: Design and performance evaluation,," *Internat. J. Robotics Res*, vol. 29, no. 5, pp. 585–597, 2010.

[6] L. Birglen, T. Laliberté, and C. Gosselin, *Underactuated Robotic Hands*. Springer, 2008.

[7] E. Brown, N. Rodenberg, J. Amend, A. Mozeika, E. Steltz, M. R. Zakin, H. Lipson, and H. M. Jaeger, "Universal robotic gripper based on the jamming of granular material," *Proceedings of the National Academy of Sciences*, vol. 107, no. 44, pp. 18 809–18 814, 2010.

[8] A. Saxena, J. Driemeyer, and A. Y. Ng, "Robotic grasping of novel objects using vision," *Int. J. Robotics Res.*, vol. 27, no. 2, pp. 157–173, 2008.

[9] E. Lopez-Damian, D. Sidobre, and R. Alami, "A grasp planner based on inertial properties," in *Robotics and Automation, 2005. ICRA 2005. Proceedings of the 2005 IEEE International Conference on*. IEEE, 2005, pp. 754–759.

[10] B. León, S. Ulbrich, R. Diankov, G. Puche, M. Przybylski, A. Morales, T. Asfour, S. Moisio, J. Bohg, J. Kuffner *et al.*, "Opengrasp: a toolkit for robot grasping simulation," in *Simulation, Modeling, and Programming for Autonomous Robots*. Springer, 2010, pp. 109–120.

[11] S. Chitta, I. Sucan, and S. Cousins, "Moveit!" *IEEE Robotics Automation Magazine*, vol. 19, no. 1, pp. 18–19, 2012.

[12] J. Bohg, A. Morales, T. Asfour, and D. Kragic, "Data-driven grasp synthesis—a survey," 2013.

[13] J. Cornelia and R. Suárez, "Determining independent grasp regions on 2d discrete objects," in *Intelligent Robots and Systems, 2005.(IROS 2005). 2005 IEEE/RSJ International Conference on*. IEEE, 2005, pp. 2941–2946.

[14] S. Ekvall and D. Kragic, "Interactive grasp learning based on human demonstration," in *Robotics and Automation (ICRA)*, vol. 4, April 2004, pp. 3519–3524 Vol.4.

[15] B. D. Argall, S. Chernova, M. Veloso, and B. Browning, "A survey of robot learning from demonstration," *Robotics and Autonomous Systems*, vol. 57, no. 5, pp. 469 – 483, 2009. [Online]. Available: http://www.sciencedirect.com/science/article/pii/S0921889008001772

[16] W. Bluethmann, R. Ambrose, M. Diftler, E. Huber, A. Fagg, M. Rosenstein, R. Platt, R. Grupen, C. Breazeal, A. Brooks *et al.*, "Building an autonomous humanoid tool user," in *Humanoid Robots, 2004 4th IEEE/RAS International Conference on*, vol. 1. IEEE, 2004, pp. 402–421.

[17] S. Niekum, S. Osentoski, G. Konidaris, and A. G. Barto, "Learning and generalization of complex tasks from unstructured demonstrations," in *Intelligent Robots and Systems (IROS), 2012 IEEE/RSJ International Conference on*. IEEE, 2012, pp. 5239–5246.

[18] R. Balasubramanian, L. Xu, P. Brook, J. Smith, and Y. Matsuoka, "Human-guided grasp measures improve grasp robustness on physical robot," in *Robotics and Automation (ICRA)*, May 2010, pp. 2294–2301.

[19] H. B. Suay, R. Toris, and S. Chernova, "A practical comparison of three robot learning from demonstration algorithm," *International Journal of Social Robotics*, vol. 4, no. 4, pp. 319–330, 2012.

[20] R. Diankov and J. Kuffner, "OpenRAVE: A planning architecture for autonomous robotics," The Robotics Institute, Pittsburgh, PA, Tech. Rep. CMU-RI-TR-08-34, July 2008.

[21] O. Zweigle, R. van de Molengraft, R. d'Andrea, and K. Häussermann, "Roboearth: Connecting robots worldwide," in *Proceedings of the 2Nd International Conference on Interaction Sciences: Information Technology, Culture and Human*, ser. ICIS '09. New York, NY, USA: ACM, 2009, pp. 184–191.

[22] H. M. Kjer and J. Wilm, "Evaluation of surface registration algorithms for pet motion correction," June 2010.

[2] ROS has an alignment function for the Barrett arm, but the sticker was positioned too close to the edge of the image to produce a reliable alignment.